



This photograph depicts one unit of the Discover supercomputer's original two-unit spinning disk metadata solution. Each unit contained 48 200-gigabyte 15K SAS drives in a RAID6 6+2 configuration, for a total usable capacity of 7.2 terabytes and a limit of ~10,000 input/output operations per second (IOPS). *Jordan Robertson, Jarrett Cohen, NASA/Goddard*




Pictured are two of the four metadata arrays installed as a part of Discover's second-generation metadata storage. Each array contained 24 500-gigabyte MLC SATA SSD drives in a RAID6 6+2 configuration, for a total usable capacity of 9 terabytes. In production, IOPS thresholds hovered around 120,000 IOPS per array. *Jordan Robertson, Jarrett Cohen, NASA/Goddard*



The photograph shows one of two recently installed Non-Volatile Memory Express (NVMe) storage arrays, now operating as the primary General Parallel File System (GPFS)-based metadata storage for Discover. These arrays increased capacity to 73 terabytes of usable space and theoretical IOPS into the millions. Initial performance testing demonstrated that each array can exceed 1 million IOPS for brief moments, with multi-day sustained IOPS of over 250,000. *Jordan Robertson, Jarrett Cohen, NASA/Goddard*

National Aeronautics and
Space Administration



IOPS Galore Encore: Upgrading a Supercomputer's Metadata with Non-Volatile Memory

The most recent wave of enhancements to the NASA Center for Climate Simulation (NCCS) Discover supercomputing cluster included an upgrade of the metadata storage. This has significantly enlarged the metadata capacity of the Discover cluster's storage, increased metadata input/output operations per second (IOPS), and introduced new technology into the General Parallel File System (GPFS) cluster.

Discover uses IBM's GPFS software to manage over 40 petabytes of user data. GPFS can store data and metadata separately, allowing system administrators to leverage high-IOPS technologies where appropriate. This year, system administrators transitioned all of the GPFS metadata over to a Non-Volatile Memory Express (NVMe) solution—a shift of over 660 million inodes (~14 terabytes across 25 filesystems).



Jordan Robertson, NASA Goddard Space Flight Center
Bruce Pfaff, NASA Goddard Space Flight Center

SUPERCOMPUTING
SCIENCE MISSION DIRECTORATE

www.nasa.gov